
A Logic Grammar Foundation for Document Representation and Document Layout

Allen L. Brown, Jr.^{†‡} and Howard A. Blair^{‡†}

[†] *Xerox Corporation*

Webster Research Center

800 Phillips Road 128-29E

Webster, New York 14580, USA

[‡] *Syracuse University*

School of Computer and Information Science

Center for Science and Technology 4-116

Syracuse, New York 13244-4100, USA

ABSTRACT: We present a powerful grammar-based paradigm for electronic document markup: coordinated definite clause translation grammars. This markup is of a declarative character, being, in effect, a collection of constraints on the logical and physical structure of documents. To the best of our knowledge, coordinated grammars and their parsers can accommodate all of the descriptive and layout processing functionality enjoyed by extant electronic markup languages. We describe an operational prototype that demonstrates the feasibility of a syntax-directed basis for formalizing and realizing document layout.

KEYWORDS: document description language, layout processing, logic grammar.

1 Introduction

Our aim is to formulate an electronic markup language with an unambiguous formal semantics within which one can specify documents in a declarative fashion. We contrast our goal with the reality of popular electronic markup languages such as $\text{\TeX}$, $\text{\LaTeX}$, Scribe, SGML and ODA. While the casual user's view of some of these markups (*e.g.* $\text{\LaTeX}$ and Scribe) would *appear* to be declarative¹, the actual meanings of user issued directives are to be understood through underlying imperative languages. This is evident when a user needs to comprehend the "style" defining mechanisms of these markups.

The technical point of view that we have adopted regarding document representation and document processing is aggressively syntax-oriented. While our

¹In complete fairness to the designers of SGML, we should point out that it is clearly their *intent* to be declarative. The problem that they leave unresolved is the formal interpretation of their declarations.

methods are related to both syntax-directed translation of programming languages and to syntax-directed natural language processing, our approach is novel in that it uses multiple grammars for the same document. Specifically, these grammars separately represent the logical and layout views of the document represented. The layout grammar is said to be *coordinated* with the logical grammar, and the two are allowed to interact through a narrowly defined interface.

In the remainder of this essay we shall demonstrate that the syntax-directed methods we have developed provide a natural and powerful framework for document representation and document processing. Our main vehicle for arriving at this conclusion is the embedding of ODA-like document representation/processing capabilities in a logic grammar framework. Presuming the reader to have some acquaintance with Prolog, we sketch a direct embedding of ODA structures and processing in that language. We introduce a particular logic grammar: the definite clause translation grammar (DCTG). We then provide a comprehensive exposition of DCTG's and parsing applied to document representation and document processing by considering the detailed specification of the layout process (realized in our operational prototype) for a simple ODA-like document. We show how to pass from the definition of the document layout process based on *total* parsing that is declarative but impractically inefficient to one based on *partial* parsing that is equally declarative and potentially quite efficient. We briefly discuss the adaptations of efficient context free parsing and incremental attribute evaluation techniques that we plan for our second phase of research. We close by sketching the future phases of research that follow from our operational prototype and its associated semantical framework.

2 ODA and Prolog

ODA [Weisz *et al.*, 1988] expresses a syntactically well-defined collection of *document constituents* of which the principal sorts are content portions (*e.g.* graphic characters, raster graphic elements and geometric graphic elements), logical objects and logical object classes, layout objects and layout object classes, and attributes. A logical (layout) structure is a tree-like arrangement of logical (layout) objects and object classes, with the trees' being "foliated" with content portion constituents.

The logical structure of an ODA document is a partitioning of the document's content based on meaning. In that context, logical object classes are elements of generic logical structure from which a set of logical objects with common characteristics may be derived (*e.g.* composite logical objects representing sections), while logical objects are elements of a document having specific

interpretations (e.g. particular chapters, sections and paragraphs).

The layout structure of an ODA document is a partitioning of the document's content based on presentation. In that context, layout object classes are elements of a generic layout structure from which a set of layout objects with common characteristics may be derived (e.g. pages with common headers and footers), while layout objects are elements of a specific layout structure of a document having specific geometric properties (e.g. particular pages and blocks). An attribute is an element of a document constituent that has a name and a value, and that expresses a characteristic of that constituent or relationship with one or more other constituents (e.g. the "presentation style" attribute establishes the relationship between a basic component description and a presentation style). Document constituent attributes can be viewed as decorating the ODA structure trees in much the same way as semantical attributes decorate parse trees in the attribute grammar paradigm.

The ODA language allows the composition of the above-mentioned constituents into *document descriptions*, each of the latter being composed of a *document profile* and a *document body*. A document profile is a collection of predefined (and preinterpreted) attributes that apply globally to the document description. The document body consists of a generic logical structure, a generic layout structure, specific logical structure, specific layout structure, and style constituents. The last are predetermined (and preinterpreted) collections of attributes that explicitly and implicitly link logical constituents with layout constituents.

We initially attempted to formalize ODA by a direct Prolog translation of ODA document constructs. We translated particular ODA document descriptions into particular Prolog fragments. Certain ODA constituents correspond to particular Prolog-defined predicates.² The real utility of ODA, however, comes only through the descriptive interpretations of various attributes (e.g. presentation style) and processes (e.g. document layout). The interpretation of these attributes is the main task of *document processing* as exemplified by the layout process. These interpretations are given in [Weisz *et al.*, 1988] in an informal fashion. The main task of formalizing ODA is to define these interpretations *rigorously*. The interpretations turn out to be other Prolog fragments relative to which we define *each* (and every) Prolog translation of an ODA document description. Hereafter, we shall refer to the translation of an ODA structural document description into a logic program as the *data* description and to the Prolog interpretation of

²In our various ODA translation efforts we have ignored the document profile and all of the details encoded in ODA's document content representations. To elucidate the core document layout process it suffices to view content portions of a document as atomic constructs with certain externally apparent attributes.

attributes (in the document processing context) as the *process* description. The latter rendering can be thought of as defining *interpreters* for various document processors.

The recipe for generating the data description is as follows: Generic (logical and layout) objects are represented as unary predicates. We may think of these generic objects as *types* whose *tokens* are specific (logical and layout) objects. In particular, tokens are individual terms in the Prolog language. Generic attributes, *i.e.* attributes of generic logical or layout objects, are represented as binary predicates. For example, the fact (part of the data description) that the generic letter has a presentation style attribute with value 'letter_layout' is represented by

```
presentation_style(X,V) :- letter(X), letter_layout(V).
```

which says that the `presentation_style` of the generic letter is the generic `letter_layout`. That a specific letter object `#Letter` had the specific presentation style `#Letter_Layout` would result from asserting the fact

```
presentation_style(#Letter,#Letter_Layout).
```

Specific (logical and layout) structures, on the other hand, are realized as Prolog terms whose nesting structure follows the hierarchy of the Prolog rules representing the related generic structures.

For the most part the translation recipe above leaves undefined the attributes that are part of the process description. For example, what does it *mean* for the logical class `summary_paragraph` to have an `alignment` attribute of `justified` (*i.e.* `alignment(X,justified) :- summary_paragraph(X).`)? The real definition for this attribute lies in its intended interpreter, in this case the *layout process*. The gist of ODA's layout model is as follows: Each ODA content portion is mapped on to one or more (layout) blocks (having geometric extents constrained by ODA layout attributes such as `alignment` having a value of `justified`) where the blocks may be generated "on the fly". The content portions are totally ordered and it is in that order that blocks for content objects are generated. The order derives from the depth-first (pre-)ordering of the tree implicit in a document's specific logical description. To capture the layout process we defined the Prolog predicate `layout_process(X,Y,U,V)` where `V` is the specific layout structure resulting from laying out the specific logical structure `U` (an instance of the generic logical structure whose user visible name is `X`) according to the generic layout structure whose user visible name is `Y`. The `layout_process` predicate can be (and has been) sufficiently Prolog-defined to handle the ODA specimen document of [Weisz *et al.*, 1988, Annex B]. We find the definition is fundamentally flawed in

that layout occurs mainly by side-effect, and its definition mimics the traditional procedural style of document layout. In the sections that follow we address these flaws by substantially raising the level of abstraction of the logic programming account of ODA data descriptions. We thereby enable an declarative account of layout process description.

3 Logic grammars and documents

Definite clause grammars are a version of context-free grammars [Harrison, 1978] that have particularly straightforward translations into definite clauses (Prolog facts and rules), yielding parsers for those grammars. There are, however, many features of languages that are either inconvenient or impossible to capture by context-free grammars. Subject-verb agreement in English is such a feature. To address the problem, investigators of logic-based parsing formalisms [Abramson and Dahl, 1989, Shieber, 1986] have borrowed liberally from researchers in attribute grammars [Deransart *et al.*, 1988]. One result of this confluence of interests is the definite clause translation grammar. This particular logic grammar provides a logic programming setting with both the context-sensitive expressive power and structuring discipline of attribute grammars. In our particular version of DCTG's (a variant of that described in [Abramson and Dahl, 1989]) we present the following context-sensitive grammar to handle noun-verb agreement:

```

sentence ::=
    meta(seq([noun_phrase^^T1,verb_phrase^^T2])) <:>
    num(Num) :- T1^^num(Num),T2^^num(Num).
noun_phrase ::=
    meta(seq([determiner^^T1,noun_phrase2^^T2])) <:>
    num(Num) :- T1^^num(Num),T2^^num(Num).
noun_phrase ::= meta(seq([noun_phrase2^^T1])) <:>
    num(Num) :- T1^^num(Num).
noun_phrase2 ::=
    meta(seq([adjective,noun_phrase2^^T2])) <:>
    num(Num) :- T2^^num(Num).
noun_phrase2 ::= meta(seq([noun^^T1])) <:>
    num(Num) :- T1^^num(Num).
verb_phrase ::= meta(seq([verb^^T1])) <:>
    num(Num) :- T1^^num(Num).
verb_phrase ::= meta(seq([verb^^T1,noun_phrase^^T1])) <:>
    num(Num) :- T1^^num(Num).
determiner ::= [the] <:> num(sing).
determiner ::= [the] <:> num(plur).
determiner ::= [a] <:> num(sing).
determiner ::= [some] <:> num(plur).

```

```

adjective ::= [decorated].
noun ::= [pieplate] <:> num(sing).
noun ::= [pieplates] <:> num(plur).
noun ::= [surprise] <:> num(sing).
noun ::= [surprises] <:> num(plur).
verb ::= [contains] <:> num(sing).
verb ::= [contain] <:> num(plur).

```

To understand the DCTG, consider the first rule of the grammar. This rule has two parts separated by the token `<:>`. The first part is a syntactic constraint indicating that a sentence is composed of a `noun_phrase` followed by a `verb_phrase`. Two Prolog variables, `T1` and `T2`, are introduced. In the course of parsing these will be bound respectively to the parse tree generated for `noun_phrase` and that generated for `verb_phrase`. The second part of the rule is zero or more (one in this case) *semantic* constraints. These semantic constraints govern the values that can be taken on by attributes associated with parse trees (and therefore with nonterminals). The parse trees associated with each of `sentence`, `noun_phrase` and `verb_phrase` have `num` attributes and the value of that attribute for a parse tree generated from `sentence` is constrained to be the same as the values of that attribute for the parse trees generated from `noun_phrase` and `verb_phrase`. The nonterminals of the grammar such as `determiner` that rewrite to terminals such as `[the]` have the values of their `num` attributes fixed at particular constants (either `sing` or `plur`). The translation (to definite clauses) of the DCTG yields *yes* on queries such as `?- sentence([some, pieplates, contain, a, surprise]).` and *no* on `?- sentence([some, pieplates, contains, a, surprise]).`

An ODA document can be embedded in the DCTG formalism in the following way: Generic logical and generic layout structures are each encoded as grammars. Nonterminals will correspond to generic objects and terminals will correspond to individual content portions. Attributes in the ODA sense will be directly mapped into attributes in the DCTG sense. Specific logical and layout structures are simply the parse trees generated by their respective grammars.

The layout structure grammar is *coordinated* with the logical structure grammar. Roughly speaking, this means that any “string” of content portions generated by the logical structure grammar is also generated by the layout structure grammar, and that certain subtrees of the parse tree of the logical structure grammar will correspond to subtrees of the parse tree of the layout structure grammar. The parse trees with respect to the two grammars for that string are distinct. The logical structure grammar for a fragment of the ODA specimen document’s [Weisz *et al.*, 1988, Annex B] generic logical structure (with some of its attributes

defined) is as follows³:

```
letter ::= meta(seq([header,body])) <:>
        object_class(root), user_visible_name("Letter").
header ::= meta(seq([date,addressee,subject,summary])) <:>
        object_class(composite), user_visible_name("Header").
summary ::= meta(rep(summary_paragraph)) <:>
        object_class(composite), user_visible_name("Summary").
```

No parse tree variables appear in this DCTG because none of the nonterminals of this fragment has attributes dependent upon the attributes of other nonterminals appearing in the same production.

4 Representing and laying out a simple ODA-like document

A full recounting of our logic grammar/parsing treatment of the data/process descriptions of the ODA specimen document would be inappropriately complex for a report of this length. Instead, we shall illustrate the essential details of our approach by appealing to an ODA representable document of considerably simpler structure. The generic logical structure of our simple document will consist of arbitrarily long sequences of paragraphs. (We shall take a paragraph to be a simple text portion.) Similarly, the generic layout structure for our simple document consists of arbitrarily long sequences of plates (pages), which in turn are arbitrarily long sequences of paragraph blocks. Below we present DCTG's `simple`⁴ and `simple_layout` that correspond to the generic logical and layout document structure diagrams illustrated in (respectively) figures 1 and 2. We begin by declaring `simple_layout` and `para_block` to be the styles corresponding respectively to `simple` and `para`:

```
styles(simple_layout,simple).
styles(para_block,para).
```

Were `simple` and `para` ODA logical objects, these declarations would correspond to asserting that the values of the "layout style" attributes of these two objects respectively have as values the generic layout objects `simple_layout` and `para_block`.

Below is the DCTG representing the generic logical structure of the `simple` document:

³Simple concatenation of phrase structures is indicated in our DCTG's by use of the metasyntactic constructor `seq`, such constructors being introduced by the indicator `meta`. We also make use of the metasyntactic constructor `rep` indicating arbitrary repetition of the phrase structure(s) in its scope. Readers familiar with ODA should note the analogy with the ODA content generator operators `SEQ` and `REP`.

⁴We identify a grammar with its root nonterminal. Hence we speak of the `simple` DCTG.




Figure 1: Simple generic logical structure.

```

simple^^T0 ::= meta(rep(para^^T1)),
  {T0^^content_interval(U,V)} <:>
  logical_type(root),
  (style(X) :- styles(X,simple)),
  (content(Z) :- sum_content_from(T1,Z)),
  (content_interval(M,N) :-
    M is 1, sum_content_from(T1,N)).
para ::= meta(seq([text^^T1])) <:>
  (style(X) :- styles(X,para)),
  (content(Z) :- T1^^content(Z)),
  (content_interval(M,N) :-
    number(M), T1^^content(U), N is (M + (U - 1))).
text ::= ["TEXT1"] <:>
  (content(Z) :- Z is 1),
  (content_interval(N,N) :- number(N)).
text ::= ["TEXT2"] <:>
  (content(Z) :- Z is 1),
  (content_interval(N,N) :- number(N)),
  layout_directive_req(apart).
text ::= ["TEXT3"] <:>
  (content(Z) :- Z is 1),
  (content_interval(N,N) :- number(N)).
  
```

The DCTG has root nonterminal *simple*, intermediate nonterminals *para* and *text*, and terminals "TEXT1", "TEXT2", and "TEXT3" (which expressions are Prolog text terms). The syntactic part of the *simple* DCTG production asserts that a *simple* is any nonempty finite sequence of *para*'s. Similarly, *para* is syntactically specified to be a *text* and *text* is syntactically specified to be one of the three terminals, "TEXT1", "TEXT2", or "TEXT3".

In addition to the syntactic characterization of nonterminals provided by productions, there is the semantic characterization provided by *guards* and attributes. The guard of the *simple* production (the expression embraced by *{}*) guarantees that the production can be used successfully only if the *content_interval*




Figure 2: Simple generic layout structure.

attribute can be given a value consisting of the ordered pair whose first and second components are the values of U and V respectively (*i.e.* the indices of the content portions spanned by `simple`). `simple` has attributes `logical-type`, `style`, `countent`, and `content.interval`. The first attribute asserts that `simple` names a logical structure grammar (*i.e.* a “root object” in the parlance of ODA). The second says that the `style` of `simple` is X if X styles `simple`, *i.e.* $X = \text{simple.layout}$. The `countent` attribute asserts that the number of content objects spanned by `simple` is the sum of the numbers of content objects spanned by each of the immediate descendants (*i.e.* the `para`’s) of `simple`. The `content.interval` attribute indicates that the interval of indices of content portions spanned by `simple` includes the first through last content portions.

The `style` attribute of `para` is X if X styles `para`, *i.e.* $X = \text{para.block}$. The `countent` attribute of `para` is the same as the `countent` attribute of the object (*i.e.* `text`) that is the immediate descendant of `para`. The `content.interval` attribute of `para` is the same as the `content.interval` attribute of the object (*i.e.* `text`) that is the immediate descendant of `para`.

Finally, the second occurrence of `text` has a layout directive request of `apart`. That is, the content associated with no previous logical object will be placed on the same layout object (a `plate` in this case) as the content associated with that occurrence of `text`. This attribution corresponds to the ODA layout directive “new layout object”. This particular content object is to be placed in a layout object distinct from that which receives the previous (in the layout order) content object. The `countent` attributes of all the occurrences of `text` have values of 1. The `content.interval` attributes of all the occurrences of `text` have values that are intervals of length 1 beginning at an index 1 greater than the upper bound of the previously indexed object (in the layout order).

Below is the DCTG representing the generic layout structure of the simple document:

```

simple_layout^^T0 ::= meta(rep(plate^^T1)),
  {T0^^content_interval(1,V)} <:>
  layout_type(root),
  (page_count(PC) ::- sum_page_count_from(T1,PC)),
  (out_trees(TTI,TTO) ::-
    styles(simple_layout,X),
    findl_node(node(X,STT,Sem),TTI,TT01,true),
    propagate_trees(T1,TT01,TTO)),
  (countent(Z) ::- sum_countent_from(T1,Z)),
  (content_interval(M,N) ::-
    number(M), sum_countent_from(T1,U),
    N is (M + (U - 1))).
plate^^T0 ::=
  meta(rep(para_block^^T1)),
  {T0^^set_depth(X),
    sum_set_depth_from(T1,S), X >= S} <:>
  layout_type(page), (page_count(PC) ::- PC is 1),
  (set_depth(X) ::- X is 1.0),
  (out_trees(TTI,TTO) ::- propagate_trees(T1,TTI,TTO)),
  (countent(Z) ::- sum_countent_from(T1,Z)),
  (content_interval(M,N) ::-
    number(M), sum_countent_from(T1,U),
    N is (M + (U - 1))),
  layout_directive_ack(apart,text).
para_block ::= meta(seq([space^^T1])) <:>
  layout_type(block),
  (out_trees(TTI,TTO) ::-
    styles(para_block,Y),
    findl_node(node(Y,STT,Sem),TTI,[T|TTO],true)),
  (countent(Z) ::- T1^^countent(Z)),
  (content_interval(M,N) ::-
    T1^^countent(U), N is (M + (U - 1))),
  (set_depth(X) ::- T1^^set_depth(X)).
space ::= ["TEXT1"] <:>
  (set_depth(X) ::- X is 0.5), (countent(Z) ::- Z is 1),
  (content_interval(N,N) ::- number(N)).
space ::= ["TEXT2"] <:>
  (set_depth(X) ::- X is 0.5), (countent(Z) ::- Z is 1),
  (content_interval(N,N) ::- number(N)).
space ::= ["TEXT3"] <:>
  (set_depth(X) ::- X is 0.5), (countent(Z) ::- Z is 1),
  (content_interval(N,N) ::- number(N)).

```

The guard of the `simple_layout` production guarantees that the production can be used successfully only if the `content_interval` attribute can be given a value consisting of the ordered pair whose first and second components are the values of `U` and `V` respectively. `simple_layout` has intermediate nonterminals `plate`, `para_block`, and `space`. The syntactic part of the `simple_layout` rule asserts that a `simple_layout` is any nonempty finite sequence of `plate`'s, that a `plate` is any nonempty finite sequence of `para_block`'s, that a `para_block` is a `space`, and that a `space` is one of the terminals "`TEXT1`", "`TEXT2`" and "`TEXT3`". Consistent with our earlier remarks that the layout structure grammar is coordinated with the logical structure grammar, the terminals of `simple_layout` subsume those of `simple`. Turning now to the semantic attributes of the `simple_layout` DCTG, we shall explain all but the `out_trees` attribute. We shall postpone its explanation until our description of the layout process.

`simple_layout` has a `layout_type` of `root`, indicating that `simple_layout` names a layout structure grammar. The value of the `page_count` attribute is constrained by its rule to be the reckoning of the number of layout objects spanned by `simple_layout` having a `layout_type` attribute with value `page`. The `content` attribute asserts that the number of content objects spanned by `simple_layout` is the sum of the numbers of content objects spanned by each of the immediate descendants (*i.e.* the `plate`'s) of `simple_layout`. The `content_interval` attribute indicates the interval of indices of content objects spanned by `simple_layout` includes the first through last (in layout order) content objects.

`plate` has a `layout_type` attribute with value `page` and a `page_count` attribute with a fixed value of unity. (Naturally, an object of `layout_type` of `page` counts as a single page!) The value of the `set_depth` attribute of an object indicates an object's vertical extent, and, in the case of a `plate`, has a value of 1.0. The `layout_directive_ack` attribute for `plate` indicates the type and the source of `layout_directive_requests` to which a `plate` is willing to respond. In this case `plate` responds to `apart` requests from (logical) objects of type `text`. As the request indicates that requesting content object is to be placed in a layout object distinct from that which received the previous content object, the acknowledgement of the request leads to the creation of a new `plate` object to receive the requesting content object. The `content` attribute asserts that the number of content objects spanned by `plate` is the sum of the numbers of content objects spanned by each of the immediate descendants (*i.e.* the `para_block`'s) of `plate`. The value of the `content_interval` attribute of `plate` is the pair consisting of the index of the first content object spanned by the left-most (in layout order) of the `plate`'s immediate descendants, and the index of the last content




Figure 3: Simple specific logical structure.

object spanned by the right-most (in layout order) of the plate's immediate descendants. The guard of the plate rule admits only those applications of the production in which the sum of the `set_depth`'s of the `para_block`'s is no larger than the `set_depth` of the plate.

A `para_block` has a `layout_type` attribute with value `block` and "synthesizes" the value of its `set_depth` attribute from the value of the same attribute of the `space` object below. A `para_block`'s `countent` attribute asserts that the number of content objects spanned by the `para_block` is the same as that spanned by its immediate descendant (*i.e.* the `space`). The value of the `content_interval` attribute of `para_block` is the pair consisting of the index of the first content object spanned (left-most in layout order) by the `para_block`'s immediate descendant, and the index of the last content object spanned (right-most in layout order) by the plate's immediate descendant.

All three `space` objects have `set_depth` attributes with values of 0.5. As indicated by the values of their `countent` attributes, each spans precisely one content object. As a consequence, their `content_interval` attributes are unit intervals whose boundaries are the indices (in layout order) of the single content objects spanned.

Considering ["TEXT1", "TEXT2", "TEXT3"] to be the input string of content portions, the logical structure grammar `simple` (figure 1) yields only one context free parse, that of figure 3. On the same input list of content portions, the layout structure grammar `simple.layout` (figure 2) yields four context free parses, figures 4-7. The main task of the layout process is to "disambiguate" the latter parses by using the context-sensitive information provided primarily by the attributes in the layout structure grammar. A collection of preference criteria is applied to the set of context free parses (of the layout structure grammar on




Figure 4: First simple specific layout structure.

a particular input string of content portions). These criteria induce a preference ordering on the parses. With respect to that ordering the “best” parses are chosen. A parse tree $P1$ is preferred to a parse tree $P2$ if

1. $P1$ satisfies the guards (the literals embraced by $\{ \}$) on all the productions used in its construction, but $P2$ does not;
2. $P1$ and $P2$ are unordered by the previous criterion, but $P1$ is coordinated with the parse of the input list according to the logical structure grammar while $P2$ is not;
3. $P1$ and $P2$ are unordered by the previous criteria, but $P1$ spans fewer page objects than does $P2$; and
4. $P1$ and $P2$ are unordered by the previous criteria, but some content portion X appears on an earlier page in $P1$ than it does in $P2$, while no content portion before (in the layout ordering) X appears on an earlier page in $P2$ than it does in $P1$.

We define a Prolog predicate `layout`

```

layout(CG,FG,L,FTT) :-
  bagof(
    FT,
    (
      parse([CG],[CT|CTT],L,[]),
      parse([FG],[FT|FTT1],L,[]),
      FT^out_trees([CT],OTT), reqs_ackd(CG,CT,FG,FT)),
    FTT2),
  min_pages(FTT2,FTT3), min_place(L,FTT3,FTT).

```




Figure 5: Second simple specific layout structure.

that guarantees an ordering under these criteria by means of other Prolog defined predicates that we shall describe presently.

`parse([CG],[CT|CTT],L,[ ])` parses the input list of content portions `L` (bound to ["TEXT1", "TEXT2", "TEXT3"]) according to the logical structure grammar `CG` (bound to `simple`), binding `CT` to the resulting parse tree. Similarly, `parse([FG],[FT|FTT1],L,[ ])` parses the input list of content portions `L` according to the layout structure grammar `FG` (bound to `simple_layout`), binding `FT` to the resulting parse tree. The success of `FT^out_trees([CT],OTT)` guarantees that the appropriate stylistic correspondences obtain between elements of the specific logical structure represented by `CT` and the specific layout structure represented by `FT` (*i.e.* they are coordinated). Recall that in the discussion above we mentioned the apart "request" and "response". The `reqs_ackd` predicate assures that for each request in the logical structure there is indeed a respondent in the layout structure. Successful acknowledgement demands (among other things) the rejection of any context-free parse that does not have the content associated with the requesting text object appearing in a `plate` distinct from that receiving the content associated with the previous text object. Among the parse trees of figures 4-7 then, only those of figures 6 and 7 are acceptable. Figure 4 is rejected by the guard of the `plate` rule and figure 5 is rejected by `reqs_ackd`. `FTT2` is now bound to a list of parse trees (bindings of `FT`) for which all the foregoing conditions obtained, that is, those of figures 6 and 7. `min_pages(FTT2,FTT3)` succeeds just in case `FTT3` is bound to those parse trees `FT` (on the list `FTT2`) having the least number of pages, (pages being those subtrees having an attribute




Figure 6: Third simple specific layout structure.

layout_type with value page). In this instance, that means exactly the parse tree of figure 6. Finally, `min_place(L,FTT3,FTT)` guarantees that `FTT` is bound to those `FT`'s on the list `FTT3` such that the content items appear as early as possible in the layout order (when compared with other members of `FTT3`) among the subtrees having an attribute `layout_type` with value page. Again, that means exactly the parse tree of figure 6.

5 Partial parsing

We have shown how the logic grammar representation of documents together with attributed parsing can give a declarative account of document layout. Our approach as described thus far would be hopelessly inefficient as a *practical* basis for document layout. In part this inefficiency is due to casting layout as an optimization problem in which we generate *all* of the candidate context-free parses, evaluate their attributes, and compare the various parses to find the optimal ones. We have remedied this particular inefficiency by resorting to *partial parsing*. We represent ordinary parse trees as Prolog terms. Each grammar generates a certain well-defined set of such terms. These trees are *total* in that they are always variable-free (ground). A partial parse tree for a grammar is (essentially) a term that unifies with some total parse tree of that grammar. We can define a preference ordering on the partial parse trees analogous to the one we defined on the total parse trees in such a way that any maximally preferred parse tree among the *total* parse trees also happens to be maximally preferred among the partial parse trees.

We can greatly restrict the parse trees that we need to consider in our




Figure 7: Fourth simple specific layout structure.

optimization problem. This follows from the fact that *layout* is the assignment of content items to pages in a manner consistent with the layout order of the content items. Thus, we are interested in the partial parses that correspond to having filled the first n pages with some initial segment of the input list of content portions. For a particular layout structure grammar and input list of content portions we define the n -page partial parse trees to be those partial parse trees generated by the grammar, each of which has n disjoint subtrees that are variable-free and whose root nodes all number among their valued attributes `layout_type(page)`, and each of which spans an initial segment of the input list of content portions. For *these* partial parse trees we define the following preference ordering: An n -page partial parse tree $P1$ is preferred to an n -page partial parse tree $P2$ if

1. $P1$ satisfies the guards on all the productions used in its construction, but $P2$ does not;
2. $P1$ and $P2$ are unordered by the previous criterion, but $P1$ is stylistically consistent (coordinated) with the portion of the logical grammar (total) parse tree spanning the same initial segment of the input list as $P1$ while $P2$ is not; and
3. $P1$ and $P2$ are unordered by the previous criteria, but $P1$ spans a longer initial segment of the input list than does $P2$.

These partial parse trees can be generated in a top-down or bottom-up fashion, and in the order of increasing n . Straightforward modification of virtually any context-free grammar parsing algorithm and associated attribute valuation algorithms will

provide a framework that will facilitate the early rejection of partial parses of which the eventually rejected total parses are substitution instances.

6 Future directions and conclusions

In order to address other sources of inefficiency in our prototype, we shall redesign and reimplement our parsers to exploit the classes of efficient context-free parsers and incremental attribute evaluators described in [Graham *et al.*, 1980] and [Reps, 1984]. Our current parsers take as input a coordinated pair of grammars and an input string of content items and produce a coordinated pair of parse trees. It is possible to compile the coordinated pair of grammars so as to generate a Prolog program specific to parsing according to those grammars. We intend to implement such a compilation strategy and thereby make additional gains in efficiency.

We have alluded to the fact that coordinated grammars (and hence document descriptions) have a declarative formal semantics. We are exploring a mathematical abstraction called a *markup scheme* [Brown and Wakayama, 1990], a formal framework modeled after program schemes [Greibach, 1975]. Markup schemes admit a least fixed point semantics derived from that of logic programs [Lloyd, 1987]. Moreover, within this framework it is possible to abstract away the details of various electronic markup languages and compare their expressive power according to the presence or absence of various features. We intend to carry out such a comparison among selected markups, examining their expressive power with respect various features of those representations.

In addition to making analytic use of markup schemes, we shall employ them in a synthetic fashion. The logic grammar formulation of document representation that we have presented is not particularly specialized to the representation of documents. We should like to conceive such a specialization, both for reasons of efficiency of document processing and to enhance the usability of the description language. To that end we hope to formulate a constraint logic programming language whose domain of discourse includes certain tree structures that describe particular documents, and whose constraints govern the admissibility of these tree structures according to structural or functional considerations.

We have offered here a powerful logic grammar-based paradigm for electronic document markup. This markup is of a declarative character, being, in effect, a collection of constraints on the logical and physical structure of documents. Moreover, this logic grammar representation admits a formal semantics that can be used directly to compare and contrast a variety of extant (and possible) electronic markup languages. To the best of our knowledge, coordinated grammars and their parsers can accommodate all of the descriptive and layout processing functionality

enjoyed by extant electronic markup languages. We have demonstrated the possibility of syntax-directed basis for formalizing and realizing document layout. We recognize that substantially more work is needed to make a syntax-directed document layout a practical reality within the coordinated grammar framework. We have embarked upon an effort to achieve that reality.

References

- [Abramson and Dahl, 1989] Harvey Abramson and Veronica Dahl. *Logic Grammars*. Springer-Verlag, New York, 1989.
- [Brown and Wakayama, 1990] Allen L. Brown, Jr. and Toshiro Wakayama. Assigning meaning to markup. Manuscript to be submitted for publication, 1990.
- [Deransart *et al.*, 1988] Pierre Deransart, Martin Jourdan, and Bernard Lorho. *Attribute Grammars: Definitions, Systems and Bibliography*, volume 323 of *Lecture Notes in Computer Science*. Springer-Verlag, New York, 1988.
- [Franchi-Zannettacci and Arnon, 1989] Paul Franchi-Zannettacci and Dennis S. Arnon. Context-sensitive semantics as a basis for processing structured documents. In Jacques André and Jean Bézivin, editors, *Proceedings of the Workshop on Object-Oriented Document Manipulation*, pages 135–146, 1989.
- [Graham *et al.*, 1980] Susan L. Graham, Michael A. Harrison, and Walter L. Ruzzo. An improved context-free recognizer. *ACM Trans. on Programming Languages and Systems*, 2(3):415–462, July 1980.
- [Greibach, 1975] Sheila A. Greibach. *Theory of Program Structures: Schemes, Semantics, Verification*. Springer-Verlag, Berlin, 1975.
- [Harrison, 1978] Michael A. Harrison. *Introduction to Formal Language Theory*. Addison-Wesley, Reading, Massachusetts, 1978.
- [Lloyd, 1987] John W. Lloyd. *Foundations of Logic Programming*. Springer-Verlag, Berlin, second edition, 1987.
- [Reps, 1984] Thomas W. Reps. *Generating Language-Based Environments*. MIT Press, Cambridge, Massachusetts, 1984.
- [Shieber, 1986] Stuart M. Shieber. *An Introduction to Unification-Based Approaches to Grammar*, volume 4 of *CSLI Lecture Notes*. Center for the Study of Language and Information, Stanford, California, 1986.
- [Weisz *et al.*, 1988] Henri C. Weisz, Ian R. Campbell-Grant, Roy Hunter, Roy Pierce, L.J. Zeckendorf, and Barry J. Woods. Information processing, text and office systems, office document architecture (ODA) and interchange format. Technical Report DIS 8613, International Standards Organization (ISO), March 1988.